Summary Comments Report

5/20/2009

Review: SciDAC Mid-Term Centers & Institutes Review 2009
Proposal Number: 100985
Proposal Title: Visualization and Analytics Center for Enabling Technologies (VACET)
Investigator: Bethel, Wes (Lawrence Berkeley National Laboratory)

1. Scientific and/or technical merit of the project

a. Does the project address an important and relevant problem facilitating breakthrough advances by the use of high performance computing in one or more science application areas?
b. What is the potential of the project to advance the state-of-the-art in HPC or HPC-enabled science?
c. Does the project demonstrate partnerships among science application scientists, applied mathematicians, and computational scientists?

Reviewer 48:

The Center has made great progress in applying visualization algorithms to the increasingly large datasets produced by the HPC community. Center personnel have developed partnerships with a large number of scientists in a variety of disciplines (climate, fusion, combustion, etc.) and have helped these scientists develop new insights into their data through visualization. I found their results compelling. I was particularly impressed by their interactions with APDEC, by the volume rendering of the high resolution Visible Man data, and by their innovations in flow visualization. Their software engineering efforts have turned VisIt into a solid, professional-quality product that has been gaining prominence within their targeted scientific communities. Overall, I view their results for the first half of their term as quite impressive.

Given their very long list of collaborators I was initially concerned that they might be spreading their resources too thinly by trying to work with too many people at once. However, they seem to have a good management structure and appear to be triaging the demands on their time fairly well.

Reviewer 37:

a. The project addresses the big data problem created by HPC applications by adapting, extending, creating when necessary, and deploying visualization and data understanding technologies to their collaborators. Their primary software is VisIt. The functionality and research is driven by their SciDAC stakeholders.

b. The knowledge from data issue is one of the most critical issues in HPC. They have one

of the only tools for big data, and they have unique features in this tool. Their innovative research, in topology, in Query-Driven Visual Data Exploration and Analysis, and other areas, is valuable to their collaborators and the wider community.

c. Yes, they have a wide range of partnerships among science application scientists, applied mathematicians, and computational scientists.

Reviewer 39:

a. Does the project address an important and relevant problem facilitating breakthrough advances by the use of high performance computing in one or more science application areas? Yes. This is directly addressed by the mission and vision of VACET. VACET's primary mission is to foster scientific insight through creating and deploying effective data understanding technology that is responsive to the needs of their stakeholders in the scientific research community who have need for such tools. The vision of VACET is to respond directly to these challenges by adapting, extending, creating and deploying visualization and data understanding technologies for their science stakeholders.

b. What is the potential of the project to advance the state-of-the-art in HPC or HPC-enabled science? As the size of the computational simulations gets larger and larger, the bottleneck in the problem solving process becomes one of managing and analyzing ever increasing stores of data. Visualization and analytics provide vital tools in this effort to extract information from this wealth of data.

c. Does the project demonstrate partnerships among science application scientists, applied mathematicians, and computational scientists? Yes, clearly it does. Section 2.1 in the report provides a fairly detailed description of the collaborators and the extent of the collaboration.

Reviewer 5:

The original project proposed addresses one of the most difficult problems in large scale high performance computing - how to effectively analyze and evaluate the often enormous output produced by the scientific simulation applications. Without advanced data management and visualization this would no longer be possible on today's petascale computing facilities and certainly pose an insurmountable roadblock for tomorrows exascale systems. This particular project proposal goes further than most and does not only focus on the visualization of the data alone, but aims for much closer integration with the scientific simulation process, data management and the analysis of the data. In this the project addresses a very relevant problem which if solved will undoubtedly lead to breakthroughs in a variety of science application areas.

Due to budget cuts the project has been reduced back in scope to visualization only, which though understandable is missed opportunity.

During the project run time it has already demonstrated its potential to help to advance the

state-of-the-art HPC enabled science through more effective data analysis. The project involves a wide range of partners from the Computational Science, Mathematics and Application Science and has demonstrated through a wide range of successful collaborations that these work effectively together in enhancing the progress of science. A worrying aspect is the near exclusive focus of this project on its own research work, little evidence was shown of collaborations with outsiders other than taken on some ideas. In particular a closer interaction with the visualization institute would seem very desirable specifically as these seem to address key basic issues of scaling visualization to the petascale range that this project is not addressing (or at least no evidence of that was produced either in the report or the presentations). The Q&A session provided some more detail and some ad hoc interaction is indeed taking place, I would strongly encourage the project to take a more formal approach to this collaboration to a) exploit the benefits of being complementary (outreach/ research/ software engineering)and b) to avoid duplication of efforts (research/ user interaction)

Reviewer 34:

The center is aimed at consolidating and coordinating efforts of leading visualization researchers to support the open science community. The development and deployment activities of related and complementary technologies such as data management, analysis, and visualization are under one umbrella allow for targeting visualization advances across multiple scientific domains. The center examines and reviews visualization technologies across is multi-organization structure and interacts with other technology centers (Mathematics and Computer Science) to address the deployment of visualization and visual analysis as part of a larger technology ecosystem. The project's key goal is to bring cutting-edge visualization and analytics to the SciDAC scientific community in the form of well-software-engineered, usable software. The application domains examined include fusion plasmas, turbulence, combustion, astrophysics, climate, environmental management, and particle accelerators.

These are relevant issues for scientific application domains that provide for breakthroughs in high performance computing. The center has already shown itself to have state-of-the-art contributions through tools like SCIRun and VisIt as well of the more recent VisTrails. These are open source tools that can be used by application scientists to visually explore complex and large data sets. These are fundamental advances to the wide application of visualization methods. The center also has a clearly demonstrated functional partnership across multiple organizations and team members that include domain scientists, applied mathematicians, computational scientists, and visualization researchers.

Reviewer 35:

a. Yes, the participants in this project are addressing application needs through software development and introduction of visualization techniques. They are providing some integration of "info vis" and scientific visualization.

b. They are increasing the ease of use of visualization tools that should increase the

productivity of HPC-enabled science. They are clearly enabling individual projects, and developing core for other areas to leverage these advances.

c. Multiple individual partnerships between scientists and scientists are documented.

Reviewer 36:

Overall, technical merit for this proposal at this stage of completion is really promising - an expansive swath in terms of breadth and in-depth synoptic views. The collaborators have successfully engaged a diverse variety of scientific disciplines to produce generalizable and scalable means of impacting investigators and simulation experimenters in their most pressing and common needs for interactive simulation and flexible visualization tools. In the case of VisIt (on Steroids), the respective coupling AND stabilization of efforts can be regarded as truly at the vanguard - since its mid-term application and adoption has been so widespread in the HPC scientific and analytical users community.

2. Appropriateness of the proposed methods or approach

a. Is the conceptual and/or mathematical framework of the enabling technology being addressed adequately developed and appropriate?
b. Does the project make use of appropriate and best available mathematical algorithms and/or computer science methods?
c. Has the project demonstrated that its software codes can effectively use the Office of Science's high performance computing resources (e.g. Leadership-class Computing Facility) or facilitate their use?
d. Are there significant current/potential problems in the project's methods or approaches? If so, is the project adequately addressing these problems?

Reviewer 48:

The center seems to have parallel, though somewhat separate efforts in visualization research and production-quality software development. Novel algorithms for doing flow visualization and volume rendering with parallel algorithms have been important research contributions. Their VisIt software is currently used on a number of supercomputers for visualization of a variety of large scientific datasets. On the whole their software and methods are excellent, though as they noted, they need to invest in implementing improved rendering algorithms.

Reviewer 37:

a. Yes, they understand the technology and address it appropriately.

b. Yes, they are leaders in the research related to this field. They use the best algorithms or develop new ones where needed.

c. Yes, the VisIt software is designed for big data. It has unique features that make

working in this environment more effective and in some cases, even possible.

d. The fact that they have been able to accomplish so much with such a large cut in the award is amazing. It would be easy to ask for more, but not really fair.

Reviewer 39:

a. Is the conceptual and/or mathematical framework of the enabling technology being addressed adequately developed and appropriate? It is clear to me that each of the individual research algorithms has a sound mathematical framework. What isn't clear to me is the mathematical framework underlying the integration of all of these techniques. It is unclear to me that there was a cohesive view of the mathematical framework for the algorithms project wide.

b. Does the project make use of appropriate and best available mathematical algorithms and/or computer science methods? Yes, the project does an excellent job of conducting research on the most challenging problems in science. All of these problems have a significant mathematical foundation. The project is staffed by a number of visualization researchers that have either an applied math or computational science background. These researchers are at the forefront of their disciplines and stay abreast of all appropriate leading edge algorithms and computer science methods.

c. Has the project demonstrated that its software codes can effectively use the Office of Science's high performance computing resources (e.g. Leadership-class Computing Facility) or facilitate their use? VisIt clearly has. It is unclear to me that the other software packages have: SCIRun, SCIRun/Fusion and VisTrails.

d. Are there significant current/potential problems in the project's methods or approaches? If so, is the project adequately addressing these problems? I am somewhat concerned that this project is not going to get the full value of all the good research that is going on in the academic community particularly with respect to transferring these algorithms and tools to the stakeholders. There are separate and in some respect competing software efforts going on within the project. I suspect that VisIt is most widely used, but there isn't quantitative evidence to back up my assumption. I know that VisTrails is somewhat of an orthogonal effort to VisIt. I am concerned that each of the individual institutions is putting their research products into their own tools rather than into a project wide software tool. I would like to see a statement to the contrary or an explanation for why this is not true.

Reviewer 5:

Unfortunately the mid term report includes little technical background details and methods are often only mentioned in passing. It is lacking in rigorous evaluations on the pro's and con's of the new technology employed in comparison to the original methodologies used and comparisons with the work of other scientists in the field.

Having said that the numerous collaborative scientific projects seem to demonstrate that

very good progress has been made and even more importantly the science application scientists and wider scientific user communities have benefited greatly from the improved technologies. The project has shown that its software can not only use high performance computing resources, but it helps to facilitate the use of even the most high end HPC facilities.

It would have been beneficial to see much more qualitative and quantitative information on the performance of the codes and the science projects addressed. In particular, in comparative analysis with other projects and to chart the progress of their own capabilities. A projection of future requirements would have been helpful too. Information on how the project will address key issues associated with petascale visualization such as I/O, communication and code scalability were only provided after much questioning in the Q&A session and did not provide a clear strategy forward. It would be beneficial to develop such a roadmap.

The project presented a wide range of user applications and stated that their users were able to use the tools effectively on their own without their support, which is very positive. Unfortunately only limited examples were presented in the report or the presentation to substantiate this. It would have been nice to see more examples of successful breakthrough science achieved by scientists without the help of VACET purely by using their codes.

Reviewer 34:

The open science community concept outlined in this proposal is conceptually sound and highly appropriate to dissemination visualization technology that examines and explores data from the DOE Office of Science's high performance computing resources. The team of researchers proposes visualization research methods that are recognized as state of the art by the ACM SIGGRAPH and IEEE Visualization communities of peer evaluators.

The key challenge for visualization research activities is insuring that domain scientists can effectively use the tools to gain insight and effectively explore massive data sets. VACET will always need to focus effort on insuring these objectives are met.

Reviewer 35:

a. The software framework is clearly being developed appropriately. However, I couldn't see an overall conceptual/mathematical framework in the research described in the report. This was addressed better in the presentation. By the end of the project the framework and contributions should be more concisely and explicitly documented so that subsequent researchers know where to pick up on ideas.
b. In partnerships, the staff are working with appropriate techniques for visualization , data processing and manipulation, and forming queries.
c. Yes, the project is demonstrating the use of codes on high performance computing resources.
d. It was not clear to me in the report how the extensions to the VisIt framework that are made for each individual project are integrated into use by future applications. This was

addressed some in the presentaitons. It is a difficult problem, but the project needs to continue to address how they can ensure that the project assets have value after the funding is complete.

Reviewer 36:

The integrated combination of prototype topologies, FTLE techniques and advanced feature/anomaly detection methods have produced quite robust AMR+ algorithms that enable 3D-4D (multi-dimensional) streamline visualizations. This has also resulted in quantitative/qualitative insights with higher potentials for anomaly detection, resolution (drill-down) and sustained speed-up in the end-to-end workflows (e.g., S3/combustion simulations, astrophysics/Chimera, APDEC AMR/VisIt parallel instantiations). The path for segmentation, scaling and benchmarking of these efforts is progressing well beyond petascale "Chi by eye" heuristics. The concurrent rigor is anchored to/with their science and collaboration stakeholders. The Morse mathematical techniques and computational geometries appear to be well wrung-out in terms of viable tool suites for examining systems' evolution and complex diffusion patterns (multi-scalar, multi-variable). Uncertainty and utility aspects are addressed and are recognized as major down-stream issues to be "float" tested with more intense feedback loops.

3. Competency of the applicant's personnel and adequacy of the proposed resources

a. How well qualified are the key personnel to carry out the research?
b. Is there an appropriate management plan for fostering coordination and collaboration among the key researchers and across the relevant disciplines and is this being implemented successfully?
c. Which SciDAC applications has the projects interacted with, and what impacts have these interactions had on the applications?
d. Even though such interactions are not specifically encouraged by the SciDAC call, has the project collaborated with other SciDAC centers and institutes (including the Outreach Center)?

Reviewer 48:

The qualifications and research accomplishments of the Center's personnel are excellent. There appears to be good communication and collaboration both within the VACET team and between the team members and individual scientific stakeholders with whom they work. A regular meeting schedule and online wiki fosters collaborative management. Their relationship with the APDEC center and with the climate researchers, among others, demonstrates close working relationships with other scientific groups.

While center personnel get regular feedback from scientific stakeholders on the specific portions of their work that relate to individual scientific communities, they do not seem to have a mechanism by which they get regular big-picture, system-wide peer review. The elimination of the external advisory board was evidently necessitated by budget cuts, but the lack of such high-level review is concerning.

Reviewer 37:

a. The key personnel are very well qualified to carry out the research. They are recognized world-wide as leaders in their fields.

b. Yes, they have a manager for each project and an executive committee that manages the flow of needs, research, and implementation choices.

c. Collaborations with Science Stakeholders

They have had visible and numerous impacts on all of their collaborations. These impacts include science impacts, parallel computing and visualization research impacts, efficiency impacts, dissemination impacts. Their collaborators are detailed below.

Accelerator
SciDAC SA: Community Petascale Project for Accelerator Science and Simulation (ComPASS), INCITE Awardee, Cameron Geddes, LBNL

Astrophysics
SciDAC SA: Computational Astrophysics Consoritium: Supernovae, Gamma Ray Bursts, and Nucleosynthesis, Stan Woosley, UC Santa Cruz and John Bell, LBN
INCITE: Multidimensional Simulations of Core Collapse Supernovae, PetaApps: Supernova Simulations with CHIMERA Tony Mezzacappa, John Blondin, Steve Bruenn

Climate
SciDAC SA: Global Cloud Resolving Model, INCITE Awardee Dave Randall, CSU
SciDAC CET: Earth Systems Grid Dean Williams, LLNL
NOAA Geophysical Fluid Dynamics Library

Combustion
Interaction of Turbulence and Chemistry in Lean Premixed Laboratory Flames, INCITE Awardee John Bell, LBNL
High-Fidelity Simulations for Clean and Efficient Combustion of Alternative Fuels, INCITE Awardee, Jacqueline Chen, SNL,

Fusion
Framework Application for Core-Edge Transport Simulations (FACETS) John Cary, Tech-X
SciDAC Center for Gyrokinetic Particle Simulation of Turbulent Transport in Burning Plasmas (P. Diamond) and SciDAC Center for Plasma Edge Simulation (C.S. Chang), Stephen Ethier, PPPL; Seung-Hoe Ku (NYU); Julian Cummings (Caltech)
Center for Extended Magnetohydrodynamic Modeling (CEMM) (S. Jardin); Center for Simulation of Plasma Microturbulence (CSPM) (W. Nevins); Simulation of Wave Interactions with Magnetohydrodynamics (SWIM), S. Kruger (Tech-X); J. Breslau (PPPL); W. Nevins (LLNL); D. Bachelor (ORNL),

Turbulence
Simulations of Turbulent Flows with Strong Shocks and Density Variations Eric Johnsen, Johann Larson (Stanford)
Reactor Thermal Hydraulics (INCITE runs of Nek5000) Paul Fischer (Argonne)

Mathematics
Applied Partial Differential Equations Center for Enabling Technology, Phil Colella, LBNL

Technology and Outreach Collaborations
Data Management
SciDAC Scientiffic Data Management Center for Enabling Technology, Arie Shoshani and Kesheng Wu, LBNL,

Outreach and Software Engineering Infrastructure
SciDAC Outreach Institute ,David Skinner, LBNL/NERSC

Code Interoperabilty
SciDAC Interoperable Technologies for Advanced Petascale Simulations (ITAPS), Lori Diachin, LLNL,

Visualization
Institute for Ultrascale Visualization, John Owens and Kwan-Liu Ma, UCD; Jian Huang, UTK

Other Collaborations

Supercomputing Center
Argonne Advanced Leadership Computing Facility, Bill Allcock (Argonne),
LBNL National Energy Research Scientific Computing Center, Katherine Yelick and Francesca Verdier (LBNL),
ORNL National Center for Computational Sciences, Doug Kothe (ORNL);

Visualization
VisIt project, Eric Brugger (LLNL)

d. Yes, they have collaborated with other centers and the Outreach Center.

Reviewer 39:

a. How well qualified are the key personnel to carry out the research? The key personnel are among the brightest minds in the visualization and analytics communities. Additionally, many of these key personnel have a substantive background in a given computational science discipline.

b. Is there an appropriate management plan for fostering coordination and collaboration among the key researchers and across the relevant disciplines and is this being implemented successfully? There doesn't seem to be anything explicitly defined to ensure coordination and collaboration among the researchers and the disciplines. That being said, there is no shortage of this happening in the project. It does not appear that the project has a formal mechanism for ensuring this, rather they rely on and leverage existing relationships with the stakeholders. I don't seem to find any mechanism for prioritizing the collaborations to ensure that high priority collaborations get the appropriate resources.

c. Which SciDAC applications has the projects interacted with, and what impacts have these interactions had on the applications? 1) accelerator, 2) astrophysics, 3) climate, 4) combustion, 5) fusion and 6) turbulence.

d. Even though such interactions are not specifically encouraged by the SciDAC call, has the project collaborated with other SciDAC centers and institutes (including the Outreach Center)? Yes, with 1) SciDAC Outreach Institute for outreach and software engineering infrastructure, 2) SciDAC Interoperable Technologies for Advanced Petascale Simulations for code interoperability, 3) Institute for Ultrascale Visualization for visualization, 4) APDEC for mathematics, and 5) SciDAC Scientific Data Management Center for Enabling Technology for data management. These are backed up by letters included at the end of the report.

Reviewer 5:

The project partners and their associates are all highly qualified, some of them with an excellent reputation world wide, this is a team well placed to deliver the project goals.

There seems to be an excellent interaction with the science application scientists and also amongst most of the researchers. It was however very difficult to extract project management information from the project beyond the conceptual, both for past and future deliverables. In particular the cycle of user need - research - development - deployment was not apparent in the management documents produced and even in discussions seemed to be quite informal. Making it difficult to assess its support for coordination and collaboration down vertically (research - development) rather than just horizontal (research - research).

The project has interacted with a range of SciDAC science application projects as well as with SDM. The collaboration with the visualization institute is rather informal and both could benefit from a more structured interaction (see previous comments.

Reviewer 34:

VACET is a multi-institutional center that includes visualization, mathematical, and computational science expertise from Lawrence Berkeley National Labs, the University of Utah's Scientific Computing & Imaging Institute, the University of California at Davis Institute for Data Analysis and Visualization, Lawrence Livermore National Laboratory,

and Oak Ridge National Laboratory.

The key personnel are recognized leaders in the field of visualization and have actively published their work in the leading computer graphics (ACM SIGGRAPH) and leading visualization (IEEE Visualization Week) conferences and journals.

The Center's functional groups are defined as: the Center PIs, the Executive Committee, the External Advisory Board, Research, and Development, Chief Software Engineer, Software Development, and Support, and Stakeholder Projects. The personnel associated with each respective functional group are experience and well qualified at addressing their appropriate management functions.

There is an effective balance between computational scientists and application domain scientists (defined as stake holders) as noted in their applications discussion.

This project is very strong on research and application, especially in regard to open source tools. It is not as strong on proactive outreach activities. While there is some mention of an annual workshop for stake holders as well as there being an exiting Wiki presence, more effort could be made to specifically target outreach activities. These outreach activities could specifically target dissemination and training in the open source tools the center provides. We would recommend some increase in outreach efforts to address this challenge and suggestion a potential collaboration with the SciDac Ultravis project that has a strong outreach effort to domain scientists.

Reviewer 35:

a. The researchers are extremely well qualified

b. There are mechanisms in place for communication. There is no group though overseeing integration in the project as a whole. If possible some involvement of an external advisory board (external to this project, not necessarily external to SciDAC) would help make sure that there really is integration between the basic research questions being pursued and the individual software engineering colloborations with projects in SciDAC science.

c. and d.
Collaborations are listed with a wide range of SciDAC applications and Centers:
SciDAC Applied Partial Differential Equation Center,
LOASIS program of LBNL,
FACETS SciDAC
COMPASS SciDAC.
Science applications:
Global Cloud Resolving model
Hybrid Numerical Model
SciDAC astrophysics consortium

Reviewer 36:

Expertise is quite impressive; high production numbers in terms of publications and presentations in collateral aspects of outreach and dissemination. Good cross-walk between government labs and higher education institutions at the right levels of interaction and communication. Ditto for load-balancing resource allocation with respect to open source community provisioning and maintenance.

4. Performance under existing award

a. Assess the progress made thus far toward the project's research goals.
b. Have the applicants disseminated the results of their research through publications in peer-reviewed journals, meetings, conferences presentations and/or other appropriate means?
c. Does the project website provide accurate, substantive and updated information about the project and its accomplishments?
d. How well is the project balancing the priorities of the applications they are interacting with and its own research?
e. Has the project effectively contributed to an efficient use of the Office of Science's high performance computing resources?

Reviewer 48:

Results have been disseminated through dozens of journal papers and conference presentations, plus VisIt tutorials. Overall the team has been quite successful at using VisIt as a vehicle for distributing their visualization developments. The website is professional and informative, but it is hard to find the VisIt wiki from the VACET page. This link should be much more prominent.

One reviewer suggested that the VisIt user base would be well served by hosting a VisIt workshop where they could discuss design decisions and receive feedback from their user group. I concur that this would be a good step to take in the future, if resources to support it could be found.

Reviewer 37:

a. The project has made outstanding progress on their research goals. Every project is generating measurable results. They are doing basic research that is decided upon as a result of customer needs and potential impact. Their research wins awards, a testament to their standing in the visualization community. The VisIt software continues to progress. It is widely used and its use is growing.

b. Yes, the applicants have disseminated their research results through numerous publications in peer-reviewed journals as well as many other venues. They have won awards for their papers and visualizations.

c. Yes, the project website provide accurate, substantive, and contains updated information about the project and its accomplishments. It is detailed in many ways.

d. The project uses its collaborations to prioritize their research. This ensures balance between their research and their collaborations

e. Yes, the project effectively contributed to an efficient use of the Office of Science's High Performance Computing resources. Their main software project, VisIt, is designed for the large data that comes out the high performance computers. It also allows interactivity and contains unique features. They have tried to make the software fit each of their collaborators needs so that their collaborators can use the software themselves. This has been so successful that it is now preferred over some previous 'home grown' software.

Also, VisIt is included with the Joule Metric, which is used for quantitatively measuring and reporting computational efficiency gains on petascale-class platforms and is the first analysis or visualization application to play this role for DOE HPC.

Reviewer 39:

a. Assess the progress made thus far toward the project's research goals. There is no doubt that this project has done a great job of furthering computational science research and visualization and analytics research as well.

b. Have the applicants disseminated the results of their research through publications in peer-reviewed journals, meetings, conferences presentations and/or other appropriate means? Yes, this project has done a tremendous job at disseminating research results through peer-reviewed journals, conferences, tutorials, and invited presentations. This is definitely a strength of this team.

c. Does the project website provide accurate, substantive and updated information about the project and its accomplishments? This is one area that I think could be improved in the project. The website does a nice job of showing off the tools that are a part of the project -- SCIRun, SCIRun/Fusion, VisIt and VisTrails. The website also does a nice job of showing off the products (images and movies) that have resulted from collaborations with scientists in the computational sciences. What the website does not do well is give much substantive information about recommendations, tips and tricks, tools for a given discipline. The website is very technology centric. Since VACET's primary mission is to foster scientific insight, I would think that having a user centric view would be more helpful. For example, under application areas and astrophysics, you could provide suggestions for which tool or tools works best and why. You could also provide tutorials and helpful suggestions or links to them. As it is now, the website presents project information and almost nothing more.

d. How well is the project balancing the priorities of the applications they are interacting with and its own research? That is a bit difficult to answer with the information given. The project is clearly advancing the state of the art in visualization. This is demonstrated by the numerous publications in peer reviewed venues. The report also spends a great deal of time presenting accomplishments with stakeholders and collaborators. What isn't pointed

out by the report is quantitative metrics for furthering the computational sciences. For example, how many publications in the computational science venues have been published as a result of the collaboration? Are any of the publications listed co-authored by stakeholders?

e. Has the project effectively contributed to an efficient use of the Office of Science's high performance computing resources? The project has clearly demonstrated a concerted effort towards porting to and running on DOE Office of Science HPC resources.

Reviewer 5:

The project has made exceptional progress in the time up to now in providing scientists with the required technology and supporting them effectively in their application and exploitation for the progress of science. Their work to date has been widely published not only in computer science specific publications, but also in some reports for the science application communities, helping with the general dissemination and uptake. There has also been a good selection of conference speeches, workshops and tutorials. For the future it would be beneficial if more of their work was published jointly in high class science journals to aid the publication and up-take of their work.

The web pages are interesting and give a very good overview of the project aims and its achievements. the resulting code is easily accessible.

The project has put a strong emphasize on its interaction with the scientists and applications they are interacting with. The reporting focuses indeed nearly to the exclusion of the more technical details on this aspect of their work. The presentation brought much more detail on the technical side, but the integration with the science application requirements was not always clear. At the beginning of the presentation a general process for capturing user requirements and feedback was laid out that should inform the research work, but in subsequent presentation this link was no longer evident.

The collaborations has certainly contributed to the more efficient usage and exploitation of HPC resources. It is however questionable at present if this model is sustainable beyond the end of this project and if the project has a long term strategy for the sustainability of their research and development outcomes.

Reviewer 34:

VACET has made significant progress with its open science focus, especially in regard to its SCIRun, VisIt and most recent VisTrails tools.

VACET has also done an effective job of disseminating it results through publications in journals and conferences, especially in regard to IEEE VIsualization Week and the annual ACM SIGGRAPH conferences. The web site is up to date and use of Wiki technology is helpful. VACET has done a good job of utilizing high performance computing resources as well as integrating applied mathematics, computer science and visualization tasks.

One recommendation pertaining to tools like VisIt would be to encourage the developers to address both online interactive high quality display as well as high quality imagery capabilities in the code development. VisIt presently does not support the development of gigapixel imagery. It is common that a visualization tool built for interactivity often does not support top of the line rendering and imagery for producing high quality prints. Since an objective is to provide leadership in visualization software codes, we would encourage the software developers to attempt to address both interactivity and high end imagery production in open science visualization tools.

Reviewer 35:

a. The development of software support for applications is clearly making good progress. There have been basic research results for many individual topics, but it is not clear how these all tie in together. It is hard to assess the progress globally as a center because the people/institutions doing each function are not identified and cross linking between the visualization projects and the scientific applications is not strongly defined in the report.

b. Many publications by project participants are listed, but it is not clear they are all related to the work of this grant. Many of the topics aren't clearly related, for example: Reeb graphs? Gene expression? Quad mesh simplification? Spline based feature curves? Acoustic metrics? A ray tracing shading language? optimal bandwidth selection for MLS surfaces? The work should be resulting in more publications in the science domains that are being supported.

c. The website provides information, but suffers the same lack of coherency as the report. As noted in the example of how to join in as a VisIt developer, it is very hard to find some of the information on the web site.

d. The balance of priorities is not clear.

e. The enhancement of VisIt to various applications has effectively contributed to the efficient use of resources.

Reviewer 36:

In my experience, the VACET presentation today was comparable to - and maybe even better than - similar program reviews on the same generic types of subject matter in analogous areas of the national security and the S&T DARPA fiefdoms (i.e.,"Bang/Byte" x "Byte/Science" value-addeds). Moreover, the DOE VACET collaboration has achieved impressive Proof-of-Concept and deployment results with half the originally intended resources. For example, the team has demonstrated leverage wrt science area applications AND the synergy of SDM's FastBit technology to filter/tune parameters based on intermediate results with much better I/O interoperability and higher, sustainable speeds.

5. Reasonableness and appropriateness of the project budget and work plan

a. Are the staffing levels and budget appropriate for carrying out the proposed research for the remaining project term?
b. Are the proposed work plans and milestones appropriate for the remainder of the project term?

Reviewer 48:

The workplan provided in response to reviewer questions is quite detailed and was evidently distilled from an even more detailed plan the team has developed internally. I believe the budget cuts are constraining the team from accomplishing some of the useful things they had planned, but given the constraints they have made reasonable adjustments. I would have liked to have seen a more concise and concrete listing of research targets and goals, with specific information regarding the size of datasets currently being used and the size of problems they anticipate addressing in the next couple years.

Reviewer 37:

a. This project was funded at a much lower level than requested. However, they have done a commendable job with the resources they have. They have a detailed plan to ensure that they will continue to succeed.

b. Yes. They have detailed milestones for each project, including research. These milestones come from interactions with the customer projects. So they are focused towards maximum impact for SciDAC.

Reviewer 39:

a. Are the staffing levels and budget appropriate for carrying out the proposed research for the remaining project term? It's difficult to know whether the staffing levels are appropriate for carrying out the proposed research because the only budget information that is provided in the report is a roll up on page 5. There are no staffing levels, just names. Additionally, I know that some of the personnel have moved institutions since the time of the original proposal and it isn't consistently reflected in the report. For example, Pascucci moved from LLNL to Utah. He is correctly identified on page 5, but is still named on the executive committee as the lead from LLNL. This being said, I know that the original budget was cut from 4M to 2.2M per year which is a very large budget cut to absorb. See below for questions regarding reduced scope.

b. Are the proposed work plans and milestones appropriate for the remainder of the project term? This is a bit difficult for me to answer. I can go back to the original proposal and look at the yearly milestones, but they are broken down by institution. I would rather have seen it presented project wide with an indicator of which institution/s is/are going to work on them. I would have like to have seen a section in the report that talked about milestones going forward. The report did a nice job of presenting successes to date, but almost nothing about future work.

Reviewer 5:

It was difficult to see what had been spent to date in contrast with the initial budget and therefore it is difficult to say if the remaining resources are sufficient to address all the objectives the project set out to tackle.

After budget cuts the project had to cut its deliverables, it would have been helpful to receive an updated set of deliverables and the rationale behind the choices made, followed by the progress made against these new deliverables. As it was it was difficult to get a complete picture. However a lot of good work seemed to have been done so far.

Unfortunately the project was unable to provide a clear and concise overview of their planned future work together with some initial timelines. Whilst they will undoubtedly deliver good work a more integrated and planned approach would be greatly beneficial to chart their achievements and help with any upcoming re-prioritizations.

Given the sparcity of quality information I feel unable to answer the questions above with confidence. I am sure that excellent work will be delivered, but am unable to say what the remaining milestones are and if they will all be achieved.

Reviewer 34:

The multi-organizational staffing levels and budgets are appropriate for the proposed research for the remaining of the project. Since the team includes representation for universities and DOE labs, there is appropriate staff to address ongoing maintenance and distribution of visualization tools via the DOE labs.

The work plans presented in regard to advanced visualization and analytics/knowledge discovery are clearly defined and targeted.

Critical to the open software ecosystem is the continued staffing level of each site having a primary software engineer whose duty it is to assure its software meets the Center software engineering criteria.

Reviewer 35:

a. The staffing levels and budgets are appropriate for performing useful research on the problems describe and supporting the collaborations listed.

b. The report doesn't explicitly tie the reported work to the timeline and deliverables in the original proposal. Clearly changes had to be made because of the change in funding level, but it wasn't clear to me from the report what the plans and milestones are for the remainder of the project term. The six page document provided at the end helps, but it is fragmented and a bit cryptic (particularly under Research Activities). It is hard to get an overall picture of what the project will achieve in the end.

To be fair to the project, it didn't seem as though the team were aware of the nature of the questions that the panelists were going to need to respond to. I think the milestones would have been more clearly and thoughtfully presented if they had been told that this was important.

Reviewer 36:

This project is an exemplar in terms of lean and mean (efficient) implementation and execution. The cohesion and synchronization of the lines of problem definition and attack are impressive. The project management oversight is well modulated and there seems to be no need for external bodies to vet the work at this point, since the project leads effectively use ex-officio, colleagues as expertise in terms of proxy advice and critique. In fact, this is to their credit (and leadership skills) in that they accepted the risk in trading off interaction with end-users to gain more traction and agile development rather than being diverted at points too premature in the development and initial deployment stages within/across the SciDAC community. The challenge will be in responding to dissemination/scale-up and push-pull with future life cycles, extended interdisciplinary use-cases, at greater scales and varied time series.

6. How well does this project advance the SciDAC goals?

Reviewer 48:

Part of the SciDAC mission is to "analyze large-scale experimental and observational data, as well as data from simulations." VACET's efforts directly address this goal by providing scientific communities with computational tools to analyze and visualize extremely large datasets. This group is one on the forefront of addressing these challenges.

Reviewer 37:

The project advances the SciDAC goals very, very well. Their research is directly targeted at discovery through advanced computation. By paying attention to the needs of their collaborators, they are even making it possible for scientists to do their own visualization using VisIt.

Reviewer 39:

VACET has delineated three measures of scientific impact: 1) provided new capabilities to science teams that aid in knowledge discovery, 2) cost savings by providing tools that have the necessary functionality rather than scientists having separate splinter efforts, and 3) improving efficiency reducing the time to solution for the scientific discovery process.

Reviewer 5:

The project says itself that there is no 'one-size-fits all' solution in this area and much

customization has to be carried out to gain the scientific impact demonstrated in the various outstanding science achievements enabled by them.

Therefore the project has certainly advanced the SciDAC goals in these cases, but it is more difficult to assess how much the wider community will benefit. This is certainly a concept which allows the achievement of far reaching scientific goals for selected scientific groups, but would require a long term support to guarantee an ongoing level of achievement comparable to what has been demonstrated so far.

Reviewer 34:

VACET clearly and specifically addresses the provision of visualization research and open science tools that support multidisciplinary projects relating to developing future energy sources, studying global climate change, accelerating research in designing new materials, improving environmental cleanup methods, and understanding high energy physics phenomena .

VACET is making a clear effort to address visualization and analytics research and development attuned to the needs of SciDAC scientists and engineers via its open science ecosystem approach. A key VACET challenge continues to be ongoing maintenance and support for their computing framework and infrastructure.

Reviewer 35:

The support for existing applications and enhancement to VisIt are excellent in making progress to use high end computing resources in scientific discovery. The research projects generally contribute to the better use of high end computing resources, but need better integration with software that delivers these capabilities to end users.

Reviewer 36:

This project advances SciDAC 2's objective to establish a salience in the composite of petascale computation, complexity and visualization arenas. The VACET team has positioned itself in multiple science applications and have demonstrated that they can meet Joule assessment criteria in truly substantive "so what/" measures. The follow-on to this seminal effort is to engage the cognition community to seed similar interactions and push/pull-like feedbacks to yield even greater rates of discovery ala induction, deduction, and abduction thinking processes. It would be interesting to see the team's "stretch" in thresholding this direction.

7. Additional Comments (including the overall strengths and weaknesses of the project):

Reviewer 48:
The highly collaborative nature of VACET's work is one of its greatest strengths, bringing visualization experts together with the scientists who can benefit from their work. Their

expertise in building high quality, efficient visualization software is their other main strength.

While I would not characterize this as a weak area, I would suggest that the team look for ways to further cultivate their user base (especially of VisIt), encourage external developers to actively contribute extensions to the code base, and solicit external input on implementation priorities. Again, I think the suggestion of a VisIt workshop is a good one. I make these suggestions as I see that VisIt had wide applicability, and the VACET team has limited resources and finite funding. Cultivating a strong open source community of both users and developers could help sustain VisIt over the coming years, especially since so many scientists are now coming to depend on VisIt for their work.

Reviewer 37:

They have done a very, very admirable job - despite having a large reduction in their award.

It would have been nice to see the analytics results, if they had been fully funded.

Reviewer 39:

VisIt is a widely used piece of software both inside and outside of the DOE research community. I am personally familiar with how well VisIt works and have been monitoring the VisIt users list. The team is extremely responsive and very helpful.

It might have been nice to see the original set of stakeholder needs, something akin to a requirements traceability matrix. I read the original proposal and did see Figure 1. I would be interested to know how you vet and handle changing user requirements. You mention that essentially all R&D activities can be traced back to stakeholder needs which are accumulated on the VACET wiki. Is there a formal process to make sure stakeholder needs are polled on a regular basis?

There was mention of the original proposal budget being cut by 45%. Was there also a reduction in scope? If so, it would be nice to see that.

There is no mention of future work and any deviations from the proposed work.

Is there a formal mechanism for coordinating and collaborating with stakeholders? Your budget was cut by 45% which means that you are running very lean. How do you prioritize these strained resources with respect to stakeholder outreach and collaboration?

Reviewer 5:

This is an outstanding project in terms of engagement and impact on the science application community, their development needs are very well researched and targeted to meet real requirements. Due to the close collaboration with the scientists a very useful

feedback loop has been created as well as a strong albeit limited user community.

The project has conceptually determined a way forward for the future sustainability of their research and development work, but will now need to work on a detailed strategy and plan to realize this concept. Following their decision to continue as open source projects after the end of the project, the entry barrier for developers who want to contribute to the projects should be lowered and more input sought.

A much stronger and more structured process is required to review and learn from competitor programs, other research efforts and user assessments of other products. Leveraging effort and input wherever possible.

Reviewer 34:

Solid proposal. VACET's provision of visualization tools that support an open science ecosystem is impressive.

There could be some improvement in more proactive outreach activities. To work towards improving outreach efforts, we would like to suggest that VACET consider a collaboration with the SciDac Ultravis project in regard to outreach efforts. We are aware that Ultravis has provided rendering modules to VACET's VisIt code (an open science visualization tool). We would like to encourage VACET to perhaps team up with UltraVis' outreach efforts and potentially encourage wider dissemination of visualization research and technologies through joint outreach efforts.

Reviewer 35:

The strength of this project is the progress made in improving VisIt to support high end computing in a variety of science domains. It is important that infrastructure is put into place so that all of this development effort will continue to be of value to the science community using high end computing after the term of the award.

The weakness in the project is the apparent gap between the definition of basic research questions addressed and the software engineering work. The team should consider as additional measures of success papers and presentations by their science colleagues that are made possible by VACET work, rather than the number of papers published in the computing literature.

Reviewer 36:

VACET team exemplifies very good unity of effort. Leadership and collegialism is distinctive. The Right Stuff. I believe they can go the distance with respect to next few years' progression along the SciDAC technical and socio-technical trajectories. When codified, the VACET work may well establish floors and ceilings on how well these methods and techniques work for various petascale-science applications and their respective renderings. A DARPA adage is that "the investment is in people, not projects."

This team is impressive. It has earned future support and continued full investment in resources from DOE/ASCR through the full term period of performance.

Overall Rating: (use whole numbers)
9-10 Excellent
7-8 Very Good
5-6 Good
3-4 Fair
0-2 Poor

Reviewer 48:
Rating: 9
Reviewer 37:
Rating: 9
Reviewer 39:
Rating: 9
Reviewer 5:
Rating: 8
Reviewer 34:
Rating: 9
Reviewer 35:
Rating: 8
Reviewer 36:
Rating: 10